

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Systems and Methods for Filter Table Optimization

Inventors:

Geary L. Eppley

Umesh Madan

David Wortendyke

ATTORNEY'S DOCKET NO. MS1-1862US

TECHNICAL FIELD

The systems and methods described herein generally relate to data processing using filters and filter tables and, more particularly, to systems and methods for optimizing filter processing by using multiple filter engines.

BACKGROUND

Computing systems - i.e. devices capable of processing electronic data such as computers, telephones, Personal Digital Assistants (PDA), etc. - communicate with other computing systems by exchanging data messages according to a communications protocol that is recognizable by the systems. A system utilizes filters to analyze messages that are sent and/or received by the system and to determine if and how the messages will be processed further.

Filter engines are used to process a diverse assortment of documents and messages against one or more filters, or queries. A filter engine may also be called an "inverse query engine." Unlike a database, wherein an input query is tried against a collection of data records, an inverse query engine inputs a data message that is tried against a collection of queries.

Most systems that utilize filter engines for processing documents and messages store multiple filters in one or more filter tables which are accessed by a filter engine when the filter engine receives an input. But the filters tables may contain thousands of entries and testing an input against each of the filters can consume a significant amount of system resources.

Filter engines developed for general use must comply with a certain language (or standard) if they are to be of use to a wide variety of applications. Compliance with a programming language requires that each and every aspect of

1 the language be supported by a filter engine, even if most of the features provided
2 in the language will never – or almost never – be used by some applications.

3 A custom filter engine may be developed for a particular application so that
4 the filter engine is streamlined to be faster and use less system overhead.

5 However, developing a custom filter engine adds a great deal of effort to
6 development of an associated application and, therefore, is usually not an option
7 preferred by developers.

8 Developers are therefore faced with a problem of either using a general
9 filter engine that supports an entire language but is inefficient, or developing a
10 custom filter engine that only support aspects of a language expected to be
11 encountered in a specific application and is therefore more efficient.

12 13 **SUMMARY**

14 At least one implementation described herein relates to filter processing
15 utilizing multiple filter engines in a way such that the use of multiple filter engines
16 is transparent to a user. A filter engine includes two or more “matchers”, each of
17 which is a filter engine in its own right. One of the matchers is a general filter
18 engine that is configured to handle all aspects of a particular language. Another
19 matcher is an optimized filter engine that handles only a subset of the particular
20 language. When an input is received, the input is analyzed to determine if an
21 optimized matcher can process it. If so, the input is directed to the optimized
22 matcher for subsequent processing. If not, the input is directed to the general
23 matcher for processing.

24 When a new filter is added to the filter engine, the filter engine analyzes the
25 new filter and determines if the filter can be supported by a grammar subset of an

1 optimized matcher. If so, the filter is added to the optimized matcher. Otherwise,
2 the filter is added to the general matcher.

3 In at least one implementation described herein, the filter engine is
4 configured to automatically supplement or generate a matcher. As inputs are
5 received, the filter engine recognizes patterns identified in the inputs. When a
6 particular pattern is seen a significant number of times, the filter engine may
7 determine that certain filters corresponding to the pattern can be isolated from a
8 general or less-optimized matcher and added to a more optimized matcher. Also,
9 the filter engine may create an additional optimized matcher to process inputs
10 matching the identified pattern. By including one or more optimized matchers,
11 many inputs may be logically evaluated against every filter in a system without
12 having to physically compare the inputs to each filter in the system.

13 14 **BRIEF DESCRIPTION OF THE DRAWINGS**

15 A more complete understanding of exemplary systems and methods
16 described herein may be had by reference to the following detailed description
17 when taken in conjunction with the accompanying drawings wherein:

18 Fig. 1 is a block diagram of a filter engine in accordance with the claimed
19 subject matter.

20 Fig. 2 is a flow diagram depicting an exemplary process to add a filter to a
21 filter engine.

22 Fig. 3 is a flow diagram depicting an exemplary filtering process to process
23 an input message.

24 Fig. 4 is a diagram of an exemplary computing environment in which the
25 implementations described herein may operate.

DETAILED DESCRIPTION

The present disclosure relates to optimizing filter table processing by utilizing a filter engine that includes multiple sub-engines, or matchers. One sub-engine is a general matcher that is fully compliant with a query language. Another sub-engine only supports a subset of the query language and, therefore, is more efficient than the general matcher. Other sub-engines may also be utilized, with each sub-engine supporting a different subset of the query language. When an input is received, the input is analyzed and directed to the optimal sub-engine that can support that particular query. Since the sub-engines are contained within a filter engine, the fact that multiple sub-engines are utilized is transparent to applications using the filter engine.

When a new filter is added to the filter engine, the filter engine is configured to analyze the new filter and determine if the filter can be associated with an optimized matcher. If the optimized matcher includes a grammar subset that supports the new filter, the new filter is associated with the optimized matcher. If the optimized matcher cannot support the new filter, then the new filter is added to the general matcher.

For discussion purposes, the terminology used herein refers to a filter engine that comprises multiple matchers, or sub-engines. The matchers represent the general and the optimized filter sub-engines discussed above. The described systems and methods refer particularly to a messaging system that includes a filter engine and one or more filter tables. The inputs received by the filter engine are messages formatted according to XML (eXtensible Markup Language), such as SOAP (Simple Object Access Protocol) messages. The filters – or queries – that

1 make up the filter tables in the exemplary systems are written in an XML query
2 language, such as XPath or XQuery.

3 Exemplary SOAP Message and XPath Query

4 The following is a representation of an exemplary SOAP message similar to
5 input messages referenced throughout this document:

```
6 <?xml version="1.0" encoding="utf-8" ?>
7 <env:Envelope xmlns:env="http://www.w3.org/2003/05/soap-
  envelope"
8 xmlns:wsa="http://schemas.xmlsoap.org/ws/2003/03/addressing
  xmlns:bk="http://booksellers.com/bookschema">
9 <env:Header>
10 <wsa:Action>http://booksellers.com/addBook</wsa:Action>
11 </env:Header>
12 <env:Body>
13 <bk:Books>
14 <bk:Book>
15 <bk:Author>Bernard Cornwell</bk:Author>
16 <bk>Title>Sharpe's Eagle</bk>Title>
17 <bk:Keyword>Military</bk:Keyword>
18 <bk:Keyword>Adventure</bk:Keyword>
19 <bk:Keyword>Wellington</bk:Keyword>
20 <bk:Keyword>Napolean</bk:Keyword>
21 <bk:Keyword>Portugal</bk:Keyword>
22 <bk:Keyword>Action</bk:Keyword>
23 <bk:PublishDate>3-13-1984</bk:PublishDate>
24 </bk:Book>
25 <bk:Book>
  <bk:Author>Stephen Pressfield</bk:Author>
  <bk>Title>The Gates of Fire</bk>Title>
  <bk:Keyword>Thermopylae</bk:Keyword>
  <bk:Keyword>Sparta</bk:Keyword>
  <bk:Keyword>Persia</bk:Keyword>
  <bk:Keyword>Xerxes</bk:Keyword>
  <bk:Keyword>Leonidas</bk:Keyword>
  <bk:Keyword>Heroic</bk:Keyword>
  <bk:PublishDate>7-26-1998</bk:PublishDate>
  </bk:Book>
</bk:Books>
</env:Body>
</env:Envelope>
```

1 The following is a representation of exemplary XPath queries similar to
2 queries referenced herein.

3 (/env:Envelope/env:Body/bk:Books/bk:Book/bk:Author = "Bernard
4 Cornwell")
5 (/env:Envelope/env:Body/bk:Books/bk:Book[bk:Author = "Bernard
6 Cornwell" or contains(bk:Author, "Pressfield")] and
7 /env:Envelope/env:Header/wsa:Action =
8 "http://booksellers.com/addBook")

9 wherein the namespace prefix mapping is:

10 env = http://www.w3.org/2003/05/soap-envelope
11 bk = http://booksellers.com/bookschema

12 The XPath language includes a number of features that a fully compliant
13 filter engine must support. Even though many of the features are rarely used, such
14 a filter engine must support them all to be compliant with the XPath language.
15 The XPath language lends itself to the present invention particularly well for this
16 reason, as well as the fact of its general popularity as an XML query language.

17 The specific examples described herein are not intended to limit the scope
18 of the claims appended hereto. The examples merely represent but one way in
19 which the claimed systems and methods may be implemented.

20 Exemplary System

21 Fig. 1 is a block diagram of an exemplary filter engine 100 in accordance
22 with the presently described systems and methods. Although particular elements
23 are shown and particular functions ascribed thereto, it is noted that a system in
24 accordance with the claimed invention may include more or less elements than
25 those shown and that functions described may be allocated among the elements in

1 a different manner than described below. The following description provides an
2 example of but one implementation of the claimed invention.

3 The exemplary filter engine 100 includes an input module 102 that is
4 configured to receive queries (i.e. filters) written in an XML query language (i.e.
5 XPath or XQuery) and messages input to the filter engine. The filter engine 100
6 also includes a sub-expression module 104 and an analyzer 106.

7 The filter engine 100 also includes three matchers, which are filter engines
8 in their own right. The three matchers are identified as a general matcher 114, an
9 enhanced matcher 116 and a fast matcher 118. The general matcher 114 supports
10 the entire query language. Each of the other matchers 116, 118 supports a unique
11 subset of the query language. In this particular example, the fast matcher 118 is
12 the most optimized matcher, i.e. the matcher that can process queries in the least
13 amount of time. The efficient matcher 116 is optimized over the general matcher
14 114, but not to the extent of the fast matcher 118.

15 The general matcher 114 includes a control module 120, a language
16 analysis module 122 and a filter table 124. The language analysis module 122
17 includes components that recognize and handle input. In the present example, the
18 language analysis module 122 includes a lexical analysis module ("lexer") 126, a
19 parser 128 and a compiler 130. The lexer 126 receives an input and groups
20 characters included in the input into elemental language units (i.e. "tokens") which
21 are then sent to the parser 128. The parser 128 analyzes the tokens received from
22 the lexer 126 and breaks them into constituent parts according to an included
23 grammar associated with the language elements supported by the general matcher
24 114. The compiler 130 translates the information from the lexer 126 and the
25

1 parser 128 into byte code or machine language that the control module 120 uses to
2 process the input against the filter table 124, which contains one or more filters.

3 The efficient matcher 116 and the fast matcher 118 include similar elements
4 included in the general matcher 114. For purposes of this description, the same
5 reference numerals are applied to the similar elements of the efficient matcher 116
6 and the fast matcher 118. Namely, each matcher 116, 118 contains a control
7 module 120, a language analysis module 122, a filter table 124, a lexer 126, a
8 parser 128 and a compiler 130.

9 As previously discussed, the efficient matcher 116 is optimized over the
10 general matcher 114. Unlike the general matcher 114, the efficient matcher 116
11 does not support an entire query language. The efficient matcher 116 supports a
12 grammar that is a subset of the query language grammar and, thus, filters 124 in
13 the efficient matcher 116 conform to the grammar subset. This means that when
14 an input message is received by the efficient matcher 116, it does not have to
15 undertake - by a significant amount - the processing load that the general matcher
16 114 would have to perform. Only input messages that can be appropriately
17 processed by filters 124 in the efficient matcher 116 are directed to the efficient
18 matcher 116.

19 The analyzer 106 is configured to identify inputs that can be handled by a
20 particular matcher by referencing the language grammar included in the parser 128
21 of the matcher. If terms included in an input are included in the grammar of an
22 optimized matcher, the analyzer 106 directs the input to the optimized matcher.
23 Conversely, if the terms are not included in the grammar of an optimized matcher,
24 the input is directed to the general matcher 114, which can handle all inputs.

1 The analyzer 106 is also configured to analyze a new filter when it is added
2 to the filter engine 100 and determine which matcher should include the new filter.
3 Similar to the process described above, the analyzer 106 determines whether the
4 grammar subset of an optimized matcher can support the new filter. If an
5 optimized matcher can support the new filter, the new filter is stored in a filter
6 table associated with the optimized matcher. Otherwise, the new filter is stored in
7 the filter table 124 of the general matcher 114.

8 Since the filter engine 100 shown in Fig. 1 includes a second optimized
9 matcher (i.e. the fast matcher 118) the analyzer 106 may also direct inputs
10 (messages and filters) to the fast matcher 118 if the inputs contain terms that are
11 included in a grammar subset associated with the fast matcher 118.

12 For purposes of this discussion, the fast matcher 118 is considered to be
13 optimized to a greater degree over the efficient matcher 116 because the grammar
14 subset of the fast matcher 118 is simpler than the grammar subset of the efficient
15 matcher 116. However, several optimized matchers may be included in a filter
16 engine without one matcher being particularly optimized to a greater degree than
17 another matcher. In such a case, neither matcher may be particularly more
18 optimized than the other.

19 Optimization

20 By way of example, consider a messaging system that transmits SOAP
21 messages and utilizes filter tables consisting of XPath queries. If there are a
22 hundred such queries in a filter table, the general matcher 114 (i.e. a general
23 purpose matcher built using an off-the-shelf XPath engine) must test each filter
24 against the input. In other words, the general matcher 114 must perform one
25 hundred tries.

1 But the filter table may contain several queries with the following
2 expression:

3 **/env:Envelope/env:Header/wsa:Action="jump"**

4 If so, it may be advantageous to include such filters in an optimized
5 matcher (i.e. the enhanced matcher 116) such as an inverse query matcher. This
6 optimized matcher would isolate a subset of the grammar to handle expressions
7 similar to the one shown above. An example of such a grammar is shown below:

```
8 EqualityExpr := AbsolutePath '=' Literal  
                | Literal '=' AbsolutePath  
9 AbsolutePath := '/' RelativePath  
10 RelativePath := Step  
                | Step '/' Step  
11 Step := NameTest  
12 NameTest := QName  
Literal := '"' [^"]* '"'  
         | "'" [^']* "'"
```

13 An optimized matcher written to handle this simple grammar could
14 be further optimized by combining individual filters into a single procedure,
15 thereby providing a substantial performance boost. One such procedure is a hash
16 function.

17 Matchers are configured to receive inputs having a certain type or structure,
18 but the inputs can have different values. A hashing function can be configured to
19 hash a particular value of an input (e.g. jump, action, remove, skip, etc.) and
20 immediately reference an entry in a hash table for the resultant hash value. As a
21 result, the input does not have to be tested against each filter in the filter table of
22 the matcher and the value must only be pulled from the input once.

23 For example, if messages differ in only one value, that value in a message
24 may be hashed. If the value were "jump", then the hash table would direct control
25

1 to the filter testing for the value "jump." As a result, filters testing for values of
2 "action", "remove", "skip", etc. would not have to be tried.

3 In the same exemplary messaging system, it may be rare for the filter table
4 to include filters with the following expression:

5 **not(boolean(//* = "Foo")) and starts-with(/env:Envelope/env:**
6 **Header/Action, "http://xyz") or /descendant::Body/..**
7 **/ancestor::Envelope/env:Header/following-sibling::Body[. = 32.333]**

8
9 Such a filter would be included with the general matcher 114 instead of an
10 optimized filter because the chances of an input satisfying the filter would be low.
11 Including the filter in an optimized matcher would increase the complexity of the
12 optimized matcher which would result in slowing the optimized matcher for the
13 cases it is configured to support. Also, if such a filter were included with an
14 optimized matcher, the filter would have to be matched against every input
15 message, which would further reduce the efficiency of the matching process.

16 Automatic Generation of Optimized Matchers

17 Optimized matchers can be automatically generated by a filter engine in
18 some implementations. In such an implementation, the analyzer 106 may be
19 further configured to keep track of inputs that the filter engine 100 receives and to
20 identify patterns that are detected more frequently than others.

21 When an input pattern is detected a significant number of times - either
22 relative to an absolute threshold or as a percentage of total inputs received - the
23 analyzer 106 can be configured to isolate filters that tend to match the dominant
24 input pattern. A restricted grammar (i.e. a grammar subset) sufficient to handle the
25 newly identified input pattern would be included in an optimized matcher.

For example, an analyzer may be configured to track similarity between inputs received by the filter engine. If the analyzer determines that, say, ten percent (10%) of the inputs it receives conform to a particular structure, the analyzer can set up a new or existing optimized matcher having filters to handle inputs conforming to that particular structure. Subsequent inputs having that structure would then be directed to that optimized matcher. If, in an alternative implementation, an absolute value were used, the analyzer could be configured to isolate the filters within an optimized matcher if, for example, it receives one thousand (1000) inputs of a similar structure within a specified period of time.

Other similar techniques may be implemented. In an alternative semi-automatic implementation, the analyzer 106 may track input patterns over time and record data pertaining to frequently detected patterns in a log. Based on the contents of the log, a developer may determine that an optimized matcher can be modified to handle the frequently detected patterns or that a new optimized matcher can be implemented to do so.

Gradual Roll-in of Optimized Matchers

One advantage realized by the systems and methods described herein is that a system utilizing inverse query engines may initially be implemented with only a general, off-the-shelf filter engine (i.e. a general matcher). As optimized matchers are developed, they can be rolled into the filter engine. Systems utilizing such a filter engine can thus be quickly brought online and gradually optimized as time and resources allow. Or inputs that may be optimized may only be discovered after a system has been operational for a period of time, thus requiring some time between implementation of a general matcher and implementation of an optimized matcher.

Sub-Expression Optimization

In addition to optimized matchers that handle entire inputs, the filter engine 100 can be configured to handle elements of inputs, i.e. sub-expressions, separately and to optimize processing of one or more sub-expressions of an input if possible.

Complex expressions in many languages, especially XPath, are composed of multiple sub-expressions. Instead of representing several smaller expressions, combining the expressions into a single expression is sometimes desired.

For example, consider the following two XPath expressions:

(1) `(/Book/Author/Name = "Melville")`

(2) `starts-with(/Book/Author, "Mel")`

Expressions (1) and (2) can also be expressed as a single expression:

(3) `(/Book/Author/Name = "Melville") and starts-with(/Book/Author, "Mel")`

Expression (3) can also be combined with a more complex XPath expression:

(4) `/descendant:Name[. = "Melville"]/ancestor::Book/*`

The complex expression that results is:

(5) `(/Book/Author/Name = "Melville") and starts-with(/Book/Author, "Mel") or /descendant:Name[. = "Melville"]/ancestor::Book/*`

Expression (5) contains features such as 'descendant' and 'ancestor'. A generalized matcher could execute the entire expression, but in a sub-optimal way.

1 Alternatively, the sub-expression module 104 (or in some implementations, the
2 general matcher 114) could decide to delegate portions of expression (5) to the
3 enhanced matcher 116 or the fast matcher 118. The techniques described herein
4 would simply be applied to a portion of each expression instead of to the entire
5 expression. For example, the sub-expression module 104 could delegate thusly:

```
6      $1 or /descendant:Name[. =  
7          "Melville"]/ancestor::Book/*
```

8 where

```
9      $1 = (/Book/Author/Name = "Melville") and starts-  
10         with(/Book/Author, "Mel").
```

11 \$1 is handled by an optimized matcher. The sub-expression module 104
12 first runs an optimized matcher 116, 118 to get the result for \$1. Then it concludes
13 by executing the remainder of the expression, if necessary. In this case, if \$1
14 returned false, the matcher wouldn't need to execute the remainder of the
15 expression.

16 It is noted that the sub-expression evaluation may only be used for sub-
17 expressions that can be executed in isolation, i.e. the result of one sub-expression
18 is not affected by the result or the action of running another sub-expression.
19 Details of sub-expression evaluation may be implemented in any one of several
20 ways not particularly discussed herein.

21 Further aspects of the elements shown and described in Fig. 1 are discussed
22 in greater detail, below, with reference to following figures.

23 **Exemplary Methodological Implementation**

24 Fig 2 is a flow diagram 200 depicting a methodological implementation of
25 an exemplary filtering process in a filter engine that includes multiple matchers.

1 In particular, Fig. 2 is directed to adding a new filter to a filter engine. In the
2 following discussion, continuing reference is made to the elements and reference
3 numerals shown in Fig. 1.

4 At block 202, a filter is received by the input module 102 to be stored in
5 one of the filter tables 124 of the filter engine 100. The analyzer 106 takes control
6 of the filter at block 204 and attempts to associate the filter with an optimized
7 matcher 116, 118 if one is available. If there is an optimized matcher available
8 that has not been tried ("Yes" branch, block 206), the analyzer 106 attempts to
9 place the filter with that optimizer. Initially, the analyzer 106 begins with the most
10 optimized matcher available - here, that is the fast matcher 118. If that matcher is
11 not appropriate for the filter, additional optimizers are tried in order according to
12 the level of optimization of the matchers.

13 At block 208, the lexer 126 breaks the filter into tokens. If that process is
14 unsuccessful ("No" branch, block 210), the matcher cannot accept the filter and
15 the process reverts to block 206. If the lexer 126 is successful ("Yes" branch,
16 block 210), the parser 128 parses the filter at block 212. If the filter includes
17 tokens that are not recognized by the parser 128, i.e. if any token is not identifiable
18 in a grammar associated with the matcher ("No" branch, block 214), the filter is
19 rejected. The analyzer 106 understands that this indicates that the matcher did not
20 accept the filter into its filter table. Therefore, the analyzer 106 must try again
21 with a different matcher at block 206.

22 If the parsing operation is successful ("Yes" branch, block 214), then the
23 compiler 130 attempts to compile the filter (block 216) into an executable form.
24 At this stage, the compiler 130 can detect operators or combinations of specific
25 tokens that are not allowed. If such a determination is made ("No" branch, block

1 218) the process reverts to block 206. If the filter is successfully compiled (“Yes”
2 branch, block 218), the filter is stored in the filter table 124 of the matcher.

3 If the filter is rejected by an optimized matcher, the next-most optimized
4 matcher is tried, if one is available. If the filter is ultimately rejected by all the
5 optimized matchers (“No” branch, block 206), then the filter is simply stored in
6 the filter table 124 of the general matcher 114 (block 220), since the general
7 matcher 114 can support any term in the language of the filter.

8 **Exemplary Methodological Implementation**

9 **Fig 3** is a flow diagram 300 depicting a methodological implementation of
10 an exemplary filtering process wherein an input is received by the filter engine for
11 processing against the filter tables 124. At block 302, an input message is received
12 by the input module 102 of the filter engine 100. The analyzer 106 analyzes the
13 input at block 304 and then tries to direct the filter to an appropriate matcher for
14 processing.

15 If an optimized matcher is not available (“No” branch, block 306), the input
16 is directed to the general matcher 114 for processing. If an optimized matcher is
17 present in the system (“Yes” branch, block 306), then the analyzer determines if
18 the optimized matcher is configured to support the particular input message. If the
19 optimized matcher can process the input message (“Yes” branch, block 308), the
20 input message is directed to the optimized filter for processing at block 310.

21 If the optimized matcher cannot process the input message (“No” branch,
22 block 308), the process reverts to block 306, where the process is repeated until an
23 appropriate optimized matcher is found or there are no optimized matchers left to
24 try. If the input message cannot be handled by any optimized matcher, it is
25 ultimately processed by the general matcher 114.

Exemplary Computer Environment

The various components and functionality described herein are implemented with a computing system. Fig. 4 shows components of typical example of such a computing system, i.e. a computer, referred by to reference numeral 400. The components shown in Fig. 4 are only examples, and are not intended to suggest any limitation as to the scope of the functionality of the invention. Furthermore, the invention is not necessarily dependent on the features shown in Fig. 4.

Generally, various different general purpose or special purpose computing system configurations can be used. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

The functionality of the computers is embodied in many cases by computer-executable instructions, such as program modules, that are executed by the computers. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Tasks might also be performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media.

1 The instructions and/or program modules are stored at different times in the
2 various computer-readable media that are either part of the computer or that can be
3 read by the computer. Programs are typically distributed, for example, on floppy
4 disks, CD-ROMs, DVD, or some form of communication media such as a
5 modulated signal. From there, they are installed or loaded into the secondary
6 memory of a computer. At execution, they are loaded at least partially into the
7 computer's primary electronic memory. The invention described herein includes
8 these and other various types of computer-readable media when such media
9 contain instructions programs, and/or modules for implementing the steps
10 described below in conjunction with a microprocessor or other data processors.
11 The invention also includes the computer itself when programmed according to
12 the methods and techniques described below.

13 For purposes of illustration, programs and other executable program
14 components such as the operating system are illustrated herein as discrete blocks,
15 although it is recognized that such programs and components reside at various
16 times in different storage components of the computer, and are executed by the
17 data processor(s) of the computer.

18 With reference to Fig. 4, the components of computer 400 may include, but
19 are not limited to, a processing unit 402, a system memory 404, and a system bus
20 406 that couples various system components including the system memory to the
21 processing unit 402. The system bus 406 may be any of several types of bus
22 structures including a memory bus or memory controller, a peripheral bus, and a
23 local bus using any of a variety of bus architectures. By way of example, and not
24 limitation, such architectures include Industry Standard Architecture (ISA) bus,
25 Micro Channel Architecture (MCA) bus, Enhanced ISA (EISAA) bus, Video

1 Electronics Standards Association (VESA) local bus, and Peripheral Component
2 Interconnect (PCI) bus also known as the Mezzanine bus.

3 Computer 400 typically includes a variety of computer-readable media.

4 Computer-readable media can be any available media that can be accessed by
5 computer 400 and includes both volatile and nonvolatile media, removable and
6 non-removable media. By way of example, and not limitation, computer-readable
7 media may comprise computer storage media and communication media.

8 "Computer storage media" includes volatile and nonvolatile, removable and non-
9 removable media implemented in any method or technology for storage of
10 information such as computer-readable instructions, data structures, program
11 modules, or other data. Computer storage media includes, but is not limited to,
12 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
13 digital versatile disks (DVD) or other optical disk storage, magnetic cassettes,
14 magnetic tape, magnetic disk storage or other magnetic storage devices, or any
15 other medium which can be used to store the desired information and which can be
16 accessed by computer 400. Communication media typically embodies computer-
17 readable instructions, data structures, program modules or other data in a
18 modulated data signal such as a carrier wave or other transport mechanism and
19 includes any information delivery media. The term "modulated data signal"
20 means a signal that has one or more of its characteristics set or changed in such a
21 manner as to encode information in the signal. By way of example, and not
22 limitation, communication media includes wired media such as a wired network or
23 direct-wired connection and wireless media such as acoustic, RF, infrared and
24 other wireless media. Combinations of any of the above should also be included
25 within the scope of computer readable media.

1 The system memory 404 includes computer storage media in the form of
2 volatile and/or nonvolatile memory such as read only memory (ROM) 408 and
3 random access memory (RAM) 410. A basic input/output system 412 (BIOS),
4 containing the basic routines that help to transfer information between elements
5 within computer 400, such as during start-up, is typically stored in ROM 408.
6 RAM 410 typically contains data and/or program modules that are immediately
7 accessible to and/or presently being operated on by processing unit 402. By way
8 of example, and not limitation, Fig. 4 illustrates operating system 414, application
9 programs 416, other program modules 418, and program data 420.

10 The computer 400 may also include other removable/non-removable,
11 volatile/nonvolatile computer storage media. By way of example only, Fig. 4
12 illustrates a hard disk drive 422 that reads from or writes to non-removable,
13 nonvolatile magnetic media, a magnetic disk drive 424 that reads from or writes to
14 a removable, nonvolatile magnetic disk 426, and an optical disk drive 428 that
15 reads from or writes to a removable, nonvolatile optical disk 430 such as a CD
16 ROM or other optical media. Other removable/non-removable,
17 volatile/nonvolatile computer storage media that can be used in the exemplary
18 operating environment include, but are not limited to, magnetic tape cassettes,
19 flash memory cards, digital versatile disks, digital video tape, solid state RAM,
20 solid state ROM, and the like. The hard disk drive 422 is typically connected to
21 the system bus 406 through a non-removable memory interface such as data media
22 interface 432, and magnetic disk drive 424 and optical disk drive 428 are typically
23 connected to the system bus 406 by a removable memory interface such as
24 interface 434.
25

1 The drives and their associated computer storage media discussed above
2 and illustrated in Fig. 4 provide storage of computer-readable instructions, data
3 structures, program modules, and other data for computer 400. In Fig. 4, for
4 example, hard disk drive 422 is illustrated as storing operating system 415,
5 application programs 417, other program modules 419, and program data 421.
6 Note that these components can either be the same as or different from operating
7 system 414, application programs 416, other program modules 418, and program
8 data 420. Operating system 415, application programs 417, other program
9 modules 419, and program data 421 are given different numbers here to illustrate
10 that, at a minimum, they are different copies. A user may enter commands and
11 information into the computer 400 through input devices such as a keyboard 436
12 and pointing device 438, commonly referred to as a mouse, trackball, or touch
13 pad. Other input devices (not shown) may include a microphone, joystick, game
14 pad, satellite dish, scanner, or the like. These and other input devices are often
15 connected to the processing unit 402 through an input/output (I/O) interface 440
16 that is coupled to the system bus, but may be connected by other interface and bus
17 structures, such as a parallel port, game port, or a universal serial bus (USB). A
18 monitor 442 or other type of display device is also connected to the system bus
19 406 via an interface, such as a video adapter 444. In addition to the monitor 442,
20 computers may also include other peripheral output devices 446 (e.g., speakers)
21 and one or more printers 448, which may be connected through the I/O interface
22 440.

23 The computer may operate in a networked environment using logical
24 connections to one or more remote computers, such as a remote computing device
25 450. The remote computing device 450 may be a personal computer, a server, a

1 router, a network PC, a peer device or other common network node, and typically
2 includes many or all of the elements described above relative to computer 400.
3 The logical connections depicted in Fig. 4 include a local area network (LAN) 452
4 and a wide area network (WAN) 454. Although the WAN 454 shown in Fig. 4 is
5 the Internet, the WAN 454 may also include other networks. Such networking
6 environments are commonplace in offices, enterprise-wide computer networks,
7 intranets, and the like.

8 When used in a LAN networking environment, the computer 400 is
9 connected to the LAN 452 through a network interface or adapter 456. When used
10 in a WAN networking environment, the computer 400 typically includes a modem
11 458 or other means for establishing communications over the Internet 454. The
12 modem 458, which may be internal or external, may be connected to the system
13 bus 406 via the I/O interface 440, or other appropriate mechanism. In a networked
14 environment, program modules depicted relative to the computer 400, or portions
15 thereof, may be stored in the remote computing device 450. By way of example,
16 and not limitation, Fig. 4 illustrates remote application programs 460 as residing
17 on remote computing device 450. It will be appreciated that the network
18 connections shown are exemplary and other means of establishing a
19 communications link between the computers may be used.

20 21 **Conclusion**

22 Although details of specific implementations and embodiments are
23 described above, such details are intended to satisfy statutory disclosure
24 obligations rather than to limit the scope of the following claims. Thus, the
25 invention as defined by the claims is not limited to the specific features described

1 above. Rather, the invention is claimed in any of its forms or modifications that
2 fall within the proper scope of the appended claims, appropriately interpreted in
3 accordance with the doctrine of equivalents.
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25